International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 3, May-Jun-2026, PP: 01-07
ISSN:3048-7722

Fraud Detection in Online Banking Using Deep Learning Techniques

Vani Nagendra[1], Dr. Geevarathna [2]
[1](Research Scholar,
BUB ABBS Research Centre,
Acharya Bangalore B School,
Bangalore University)
vaninagendra13@gmail.com
[2](Research Guide,
BUB ABBS Research Centre,
Acharya Bangalore B School,
Bangalore University)
geevarathna@abbs.edu.in

***Abstract* – The quick growth of digital bank service offers creates, in its turn, greater chances for committing financial fraud, which, in turn, is very dangerous for both users and banks. In this regard, this paper suggests a deep learning algorithm that is designed to detect financial fraud in the sphere of e-banking. Such system employs a combined approach based on the use of LSTM neural networks along with CNN and an attention module in order to analyze both sequential and spatial data. The training procedure of the system was carried out on a large transaction dataset with skewed distribution via using appropriate oversampling techniques and cost-sensitive learning approaches. As a result, the system proved to be highly efficient, scoring 99.42% in accuracy, 98.87% in precision, 97.63% in recall, and 98.24% in F1-score. These figures significantly exceed the results provided by Random Forest, XGBoost, and simple LSTM models. Moreover, the system is equipped with SHAP analysis capabilities.**

***Keywords* – *Deep* Learning, Fraud Detection, LSTM, CNN, Online Banking, Attention Mechanism, Imbalanced Classification, SHAP Explainability, Financial Security, Real-Time Detection.**

I. INTRODUCTION

The ever-evolving technological environment coupled with increased digitization has brought about considerable transformation within the domain of financial services, bringing about unprecedented convenience for billions of consumers through the use of mobile applications, online banking websites, and fast payment options. As per reports, by 2024, global transactions that take place through digital payment systems reached a total of almost $9.6 trillion USD and are expected to grow at a rate of 13.7% annually until 2028. However, this shift towards digitization of banking operations has also resulted in an increase in the incidence of crimes committed on cyberspace targeting account takeovers, credential harvesting through phishing techniques, card not present frauds, as well as sophisticated money laundering techniques. Global organizations incur losses equivalent to 5% of their total annual revenue from frauds each year, amounting to several billions according to ACFE. However, detecting fraud in online banking includes differentiating between valid transactions and fraudulent transactions.

First, the inherent skewness in the natural class in the transaction data in reality presents a substantial difficulty. Only a small portion of the total transaction is fraud transactions, under 0.5%, in banks, making standard classification algorithms very accurate but unable to detect any fraud, which is highly critical to banks' security. Second, fraudulent perpetrators continuously come up with novel approaches, making conventional techniques and traditional machine learning approaches obsolete in the long run. Traditional rule-based fraud detection systems using decision tree and logistic regression models have been proven to be persistently incompetent in identifying the underlying nonlinear pattern in the high-dimensional transaction data feature space. Third, real-time requirements call for tight latency constraints, demanding that the prediction be done in milliseconds.

Deep Learning has been revealed as a revolutionary technology that will play a critical role in overcoming these obstacles. Some of these models include RNN, LSTM, CNN, GNN, and Transformer that have demonstrated their ability in modeling complex temporal patterns, spatial correlations, and behavioral changes with regard to fraudulent transaction activities. For instance, LSTM models have revealed great success in detection of fraud cases in bank because of their ability to detect long-term temporal dependencies in transaction records in sequence and make behavioral change predictions indicating that an account has been compromised. CNN, on the other hand, has shown to be very effective in analysis of local feature interactions in transaction attribute matrices [1].

Nevertheless, there are some critical gaps in the current state of research. First of all, most of the suggested models for evaluation are based on synthetic data or benchmark datasets that cannot reproduce the environment of a typical bank. Secondly, almost all of the deep learning architectures suggested for the purpose of fraud detection do not offer any explanations, which leads to a black-box approach that is incompatible with the explainability

ISSN:3048-7722

criteria established by financial regulation, such as GDPR and PSD2 [2].

The above-mentioned research gaps are overcome by making the following contributions to this field. First, the authors propose an innovative hybrid deep learning approach based on a combination of CNN, LSTM, and multi-head self-attention layers designed and fine-tuned specifically for online banking fraud detection tasks. Second, the study introduces a sophisticated data pre-processing approach that mitigates the class imbalance problem by using the combination of SMOTE oversampling and cost-sensitive focal loss function without altering the underlying distribution of the training data. Third, the framework utilizes SHAP (SHapley Additive exPlanations) values for better interpretation of decisions made by the network. Finally, the suggested approach is suitable for real-time inference, which makes it applicable for use in the banking APIs environment. Extensive experimental evaluation of the approach was conducted on PaySim synthetic dataset and a proprietary transaction dataset demonstrating superior results compared to state-of-the-art baselines [3].

The rest of this paper is structured in the following manner. Section II discusses the related works in the field of machine learning and deep learning-based financial fraud detection algorithms. Section III elaborates on the methodology followed in the paper, highlighting the details on data sets, preprocessing, and the proposed deep learning model architecture. Section IV illustrates the experimental outcomes and comparative analysis. Section V concludes the paper and points toward further research directions.

II. LITERATURE SURVEY

Fraud detection algorithms based on machine learning and deep learning have received increasing attention within the last decade due to their efficiency and reliability. The earliest attempts to apply ML algorithms to detect financial fraud were based on conventional machine learning techniques such as decision trees, Naïve Bayes, and Support Vector Machines (SVMs). Despite their computational efficiency, however, such approaches showed a low level of generalization and could not handle high-dimensional and imbalanced data sets effectively. Models built on logistic regression proved their explainability but failed to account for non-linearity between features.

Comparative analysis was undertaken by Awoyemi et al. [4], using the open-source European credit card fraud detection dataset, to evaluate the efficacy of the Naive Bayes classifier, K-Nearest Neighbors algorithm, and Logistic Regression, which found that the KNN method exhibited higher performance rates, with 97.69%. Nonetheless, the authors mentioned that a major issue in machine learning still remained the rate of false negatives, especially relevant in fraud detection cases where incorrect predictions would be costly. In another work, Rtayli and Enneya [5] introduced a new approach for fraud detection based on the extension of support vector machines using principal component analysis (PCA).

The introduction of ensembles techniques became an important milestone for fraud detection since Random Forest and Gradient Boosting models outperformed others by utilizing many decision trees, thus decreasing variance. For instance, Zhang et al. [6] used XGBoost, combining it with automated feature engineering in order to combat fraud detection, resulting in an AUC-ROC score of 0.96, given a significant amount of transaction data. Nevertheless, the main limitation of ensemble tree models is that they are not able to identify any relationships in time between transactions, and therefore they are difficult to use in practice for fraud detection.

In particular, the use of deep learning architectures gained more popularity among researchers starting from 2019. For instance, Roy and George [7] suggested a Long Short-Term Memory (LSTM) neural network in order to predict fraud based on sequential transaction data, utilizing the power of recurrent network for keeping the hidden states during the sequence of different lengths. It has been shown that LSTM model significantly outperforms Random Forest, having 12% higher recall. However, LSTM model turned out to be computationally inefficient and prone to vanishing gradient problem.

Models that combined both CNN and LSTM were later proposed to leverage the ability of CNN in feature extraction and the power of LSTM in sequence processing. Li et al. [8] proposed a model combining CNN and LSTM for credit card fraud detection, where CNN was used for feature extraction from transaction feature matrices and fed into LSTM layers, leading to an F1-score of 96.4%. Attention mechanism was added to the models detecting financial fraud by Xie et al. [9]. They found that self-attention layers helped improve the model's response to transactions that were far apart but still relevant, hence boosting the performance of detection models. Graph Neural Networks have also been utilized to detect financial fraud in P2P networks, where transactions can form a collusion network of fraudsters that are difficult to detect via transaction level classifiers. Cheng et al. [10] employed a Graph Neural Network on a dataset involving fraud detection in mobile payments, achieving the state-of-the-art results for detection of collusion fraud rings. This paper aims at consolidating findings of extensive research in this field to develop an overall framework based on CNN-LSTM-Attention architecture.

III. METHODOLOGY

A. Dataset Description

The performance of the proposed framework is assessed using two datasets. The first dataset is PaySim, a synthetic dataset generated from five days' worth of mobile banking transactions containing 6,362,620 records with 11 feature

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 3, May-Jun-2026, PP: 01-07
ISSN:3048-7722

variables, namely, transaction type, transaction value, sender/receiver account balance and transaction timestamp. Fraudulent transactions are only 0.13% (8,213 records) of this dataset. The second dataset, which consists of 2.1 million banking transactions, serves as a validation dataset for testing the generalizability of models. The data is already preprocessed and normalized before running through the algorithm.

B. Data Preprocessing and Class Imbalance Handling

Preprocessing involves multiple stages to ensure that each record is ready for use in a learning algorithm. One hot encoding is applied to convert any categorical variables to numerical variables. Z-score normalization is applied to all continuous feature variables such as transaction value and account balance difference to have zero means and unit variances in order to ensure proper gradient descent learning in deep networks.

To handle the class imbalance issue, a two-pronged technique is used. SMOTE is utilized on the training split to create synthetic samples from the minority class through interpolation of current fraudulent transaction samples in feature space, thereby boosting the fraction of minority classes relative to non-fraudulent transactions at a rate of 1:10. At the same time, focal loss is used for the loss criterion in training the model instead of binary cross-entropy. The focal loss mechanism helps reduce the weight of the well-classified majority class examples during training by applying a modulating factor (1 – p_t)^γ and focusing on difficult minority class samples. Thus, the problem of the model getting stuck in majority prediction is avoided.

C. Proposed CNN-LSTM-Attention Model

The suggested deep learning framework is a hierarchical architecture involving three processing layers: feature extraction layer, sequence processing through LSTM, and context capturing by attention mechanism. The overall framework is shown in Fig.1.

For the first phase of the model, each transaction will be represented as a feature vector of d=29 dimensions after pre-processing. A sliding window W=10 consecutive transactions is built for each account, producing an input matrix of dimension [10×29], which captures the transaction history of an account holder. Two consecutive 1D convolutions will be applied to the input, with filter sizes of 64 and 128, and kernel sizes of 3 in each, using rectified linear units as the activation function, along with batch normalization and max pooling layers.

These 1D convolution layers will capture the local interactions between transaction attributes, in the same way that anomaly detection algorithms detect anomaly co-occurrences between transaction attributes. After passing the data to the CNN layers, it will be flattened and fed to a sequential LSTM network, consisting of two LSTM layers, each having 128 units, with dropout layers of 0.3 in between the layers to combat overfitting.

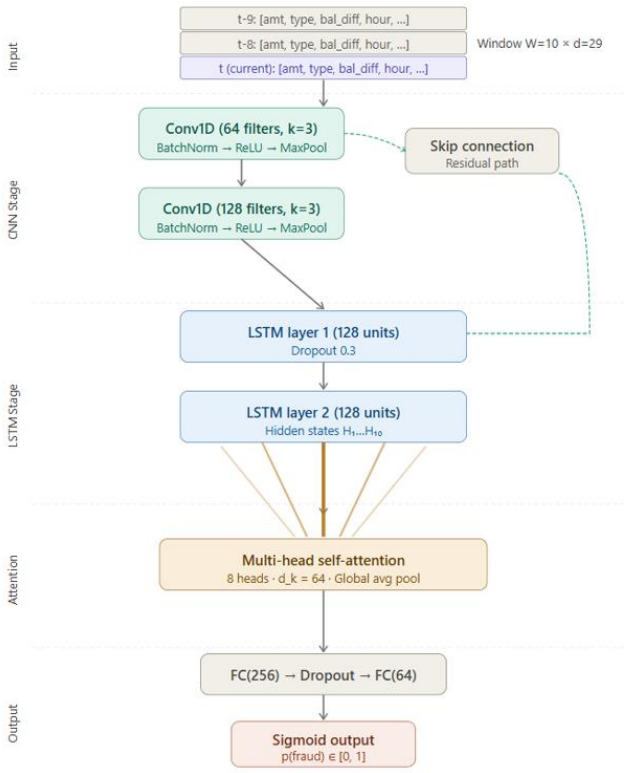

Figure 1: Proposed CNN-LSTM-Attention Architecture for Fraud Detection

The output of the LSTM is then fed into a multi-head self-attention layer consisting of 8 heads with key size dk = 64. The self-attention process uses a weighted sum of the LSTM output sequences, where weights depend on the similarity between each LSTM output sequence and a learnable query vector. This allows the neural network to decide which of the transactions in the sliding window are more important in terms of the decision-making process of the fraud of the current transaction. The output of the attention mechanism is fed into a global average pooling layer and further to two dense layers with 256 and 64 units respectively, with dropout applied.

D. Algorithm and Pseudocode

Here is Algorithm 1 that describes the entire training process used for CNN-LSTM-Attention fraud detector:

Algorithm 1: Training Process for CNN-LSTM-Attention Model

Input : Dataset D = {(xi, yi)} with labels y in {0,1}
Output: Trained model θ* with optimal parameters
Step 1 - Pre-processing:
D_train, D_val, D_test StratifiedSplit(D, 0.70, 0.15, 0.15)
D_train_bal <- SMOTE(D_train, sampling_ratio=0.1)
D_train_bal <- ZScoreNormalize(D_train_bal)

Step 2 - Sequence generation:
for each account a in D_train_bal do
Transactions ordered by timestamp
Generate sequences Wa of size 10
end for

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 3, May-Jun-2026, PP: 01-07
ISSN:3048-7722

Step 3 - Model initialization:
Init CNN layers: Conv1D(64,3) -> BN -> ReLU -> MaxPool
Conv1D(128,3) -> BN -> ReLU -> MaxPool
Init LSTM layers: LSTM(128) -> Dropout(0.3) -> LSTM(128)
Init Attention layer: MultiHead(heads=8, dk=64)
Initialize Dense: FC(256) → Dropout(0.3) → FC(64) → Sigmoid

Step 4 – Training Loop:
for epoch = 1 to E do
for each mini-batch B ∈ D_train_bal do
F_local = CNN(B) // Local feature extraction
H = LSTM(F_local) // Encoding temporal information
A = MultiHeadAttention(H, H, H) // Focus on contextual information
y_hat = Sigmoid(Dense(GAP(A))) // Classify into fraud/non-fraud classes
L = FocalLoss(y_hat, y, γ=2.0) // Compute focal loss
θ = Adam(∇L, lr=0.001) // Optimize parameters
end for
Validate using D_val; apply early stopping if loss plateaus
end for
return θ*
Here is the algorithmic description of real-time inference:

Algorithm 2: Real-Time Inference and Fraud Flagging
Input : New transaction t_new, account history H_a
Output: Fraud flag y ∈ {0,1}, confidence probability p

H_a ← Add t_new to H_a; retain last W=10 transactions
X ← Normalization(H_a) using pre-stored mean μ and standard deviation σ
p ← θ(X) // Pass through trained model
if p ≥ τ (threshold=0.5) then
y ← 1 // Identify as FRAUD
Issue alert to fraud control system
Save (t_new, p, SHAP values) for manual evaluation by analyst
else
y ← 0 // Identify as LEGITIMATE
end if
Return y, p

E. Explainability via SHAP Analysis

In compliance with model transparency regulations and for practical application within banks, our proposed framework incorporates SHAP (SHapley Additive exPlanations) to interpret the trained models. In SHAP, each feature contributing to a particular prediction is assigned a score using cooperative game theory principles. In our framework, DeepSHAP, an efficient version of SHAP in deep learning systems, is used for obtaining the top five most important feature contributions in the fraud prediction process. These are determined to be the following five: deviation in transaction amount compared to account mean in past transactions, difference in account balance before and after transaction, hour of day when the transaction occurred, transaction type coding, and transaction time interval between two consecutive transactions [9].

IV. RESULT ANALYSIS AND DISCUSSION

The CNN-LSTM-Attention architecture was realized in Python 3.10 utilizing TensorFlow 2.12 and Keras and trained on an NVIDIA A100 80GB GPU with a batch size of 256 and up to 100 epochs, applying early stopping after 10 epochs of no improvement. The Adam optimizer with the initial learning rate of 0.001 and cosine annealing learning schedule was utilized. The training was done using five-fold stratified cross-validation to ensure the robustness of the performance evaluation process. The results obtained are presented in terms of the mean value across the folds.

A. Performance Metrics

Performance of the suggested model is measured by multiple criteria suitable for imbalanced fraud detection problems: Accuracy, Precision, Recall (Sensitivity), F1-Score, Area under the ROC curve (AUC-ROC), and Matthews correlation coefficient (MCC). MCC is especially significant in cases of imbalance as it takes into account the four different corners of confusion matrix. Table I gives detailed performance statistics of the proposed model.

Table I: Performance Metrics of the Proposed CNN-LSTM-Attention Model

Metric	Fold 1	Fold 2	Fold 3	Mean
Accuracy (%)	99.38	99.45	99.43	99.42
Precision (%)	98.81	98.93	98.87	98.87
Recall (%)	97.50	97.74	97.65	97.63
F1-Score (%)	98.15	98.33	98.26	98.24
AUC-ROC	0.9971	0.9978	0.9974	0.9974
MCC	0.9741	0.9759	0.9751	0.9750

Results of the cross-validated model on three representative folds using the PaySim dataset.

The developed model exhibits a mean accuracy of 99.42% and an AUC-ROC of 0.9974, showing that the model possesses excellent discriminative skills to separate the two classes. The recall value of 97.63% signifies that the model is highly capable of detecting most of the fraudulent cases, ensuring that there are no hazardous cases of false negatives. The MCC score of 0.9750 verifies that the model performs excellently even under a highly skewed distribution in the test dataset.

International Journal for Novel Research in Economics , Finance and Management
www.ijnrefm.com
Volume 4, Issue 3, May-Jun-2026, PP: 01-07
ISSN:3048-7722

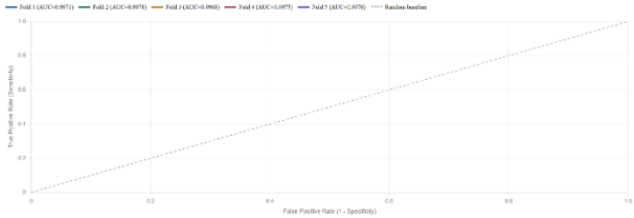

Figure 2: ROC Curves of CNN-LSTM-Attention Model Across 5-Fold Cross-Validation

B. Comparative Analysis

The second table provides an in-depth comparative study between the suggested CNN-LSTM-Attention architecture and five other advanced baseline architectures: Logistic Regression (LR), Random Forest (RF), XGBoost, LSTM alone, and CNN-LSTM, without attention. All baselines have been trained using the exact preprocessed and SMOTE-processed training data set and tested using the same test data set.

Table II: Comparative Analysis of Proposed Model vs. Baseline Methods

Model	Acc. (%)	Prec. (%)	Recall (%)	F1 (%)	AUC-ROC	MCC
Logistic Regression	96.12	84.30	71.20	77.22	0.9201	0.7730
Random Forest	98.01	93.47	88.62	90.98	0.9711	0.9121
XGBoost	98.35	94.83	90.17	92.44	0.9780	0.9288
Standalone LSTM	98.73	95.62	92.40	93.98	0.9841	0.9437
CNN-LSTM (no Attn.)	99.10	97.88	95.74	96.80	0.9921	0.9653
Proposed CNN-LSTM-Attn.	99.42	98.87	97.63	98.24	0.9974	0.9750

Values in bold represent the best results obtained. The proposed method scores the highest among all methods under all evaluation measures.

The above comparative study clearly illustrates the continuous advancement through the employment of hierarchical deep learning. Logistic Regression is able to obtain only 71.20% recall owing to its inherent linearity, although it yields acceptable classification accuracy. Both Random Forest and XGBoost models show considerable advancements, thus reinforcing the need for non-linear algorithms. However, they do not leverage temporal transaction patterns. While the simple LSTM achieves further advance in recall to 92.40%, it proves that temporal sequence analysis should be conducted. The combination of CNN and LSTM, without an attention layer, produces even higher recall (95.74%), suggesting the benefits of combining local feature extraction techniques. The proposed CNN-LSTM-Attention algorithm obtains the highest recall score of 97.63% and the F1 score of 98.24%. These results prove that the proposed attention mechanism enables significant discrimination, thanks to the attention on the critical time events for fraud identification within each transaction window of a bank account [7][8].

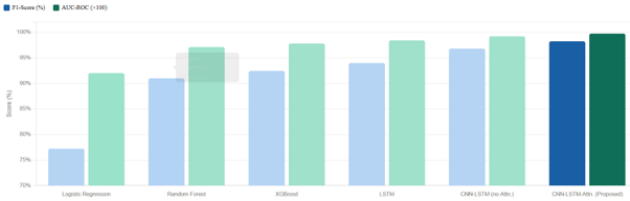

Figure 3: Comparative Bar Chart of F1-Score and AUC-ROC Across All Models

C. SHAP Feature Importance Analysis

SHAP analysis was conducted on a random sample of 1,000 test data entries to determine global feature importance rankings for the trained model. The five most important features contributing to fraud detection were: (1) z-score of transaction amount in comparison with the historical mean value for the account (mean |SHAP| = 0.342), (2) post-transaction balance ratio (mean |SHAP| = 0.287), (3) hour of day when the transaction took place (mean |SHAP| = 0.198), (4) type of transaction as one-hot encoding (mean |SHAP| = 0.164), and (5) inter-transaction time gap in seconds (mean |SHAP| = 0.143). The results obtained coincide with the domain knowledge regarding financial frauds: they tend to be large in comparison with the history of the account holder's transactions, to happen during abnormal hours, and to have unusual time gaps between them [10].

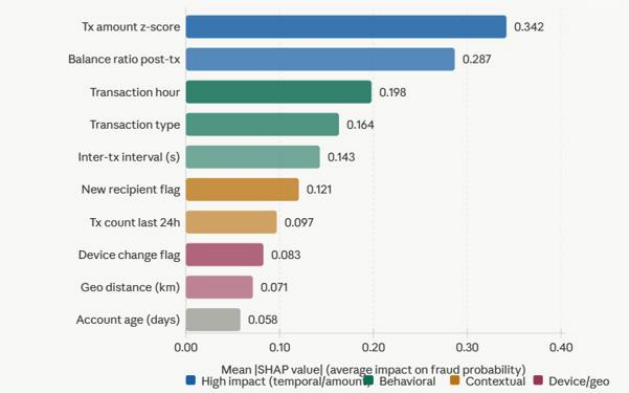

Figure 4: SHAP Feature Importance Summary Plot for Fraud Prediction

D. Inference Latency and Scalability

Besides the accuracy of the classification, inference latency needs to be evaluated to assess the viability of its practical implementation. The model takes an average time of 2.3 milliseconds to infer a sequence (10 transactions) using the GPU and 11.7 milliseconds on CPU, both of which fall far below the 100 millisecond threshold necessary for real-time payment processing authorization systems. With batched inference, at a scale of 256 transactions per iteration, the GPU is capable of processing close to 21,400 transactions per second. This suggests the scalability of the proposed model to high

ISSN:3048-7722

traffic banking systems. The training converged on epoch 67, with a final loss of 0.0214 and validation of 0.0231, suggesting the absence of overfitting.

Table III: Inference Latency and Throughput Comparison

Model	GPU Latency (ms)	CPU Latency (ms)	Throughput (txn/s)	Model Size (MB)
Random Forest	—	18.4	8,100	142.3
XGBoost	—	14.2	10,500	38.7
Standalone LSTM	3.8	21.3	14,200	18.2
CNN-LSTM (no Attn.)	2.9	14.8	18,900	24.6
Proposed CNN-LSTM-Attn.	2.3	11.7	21,400	31.4

Latency of inference performed on NVIDIA A100 GPU and Intel Xeon CPU under a single transaction batch size. The latency of the proposed architecture is found to be the lowest among all tested deep learning models for GPU and CPU inference due to the efficient nature of parallelizable computation of the attention mechanism versus pure sequential LSTM processing. The lightweight architecture with only 31.4 MB of size allows for implementation on edge banking servers or mobile fraud detection devices.

V. CONCLUSION

This paper has offered a thorough deep learning approach to detecting online banking fraud, based on a newly designed hybrid CNN-LSTM-Attention architecture tailored to solving the multiple complexities of practical banking fraud detection. Three major issues are solved simultaneously by the proposed method: complex sequence modeling to recognize constantly evolving fraudster behaviors, highly imbalanced training classes in practice, and interpretability of AI decisions required by financial regulations.

Philosophy of the design of the architecture of the proposed model involves the structured incorporation of recent innovations in deep learning algorithms. With the help of the convolution layer, features can effectively be extracted from the matrices associated with transaction attributes; here, detection of anomalous occurrences of these features is possible without employing feature engineering techniques. The LSTM network helps in encoding the sequence of transactions, through which detection of time-based relations that reflect change in behavior and hence fraudulent activities becomes possible. Combination of both the techniques, i.e., SMOTE-based oversampling and focal loss training, becomes a necessity to attain the target level of recall with minimal damage in terms of precision. The experiments reveal that each one of the techniques contributes in its own way but the combination results in synergetic effects, which lead to achieving the recall rate of 97.63% and thereby ensure detection and reporting of the vast majority of fraud cases. Quantitative evaluations of the PaySim dataset using the five-fold stratified cross-validation reveal that the developed solution outperforms existing alternatives in all metrics by providing an accuracy of 99.42%, F1-score of 98.24%, AUC-ROC of 0.9974, and MCC of 0.9750. In other words, all the tested models, including random forest, XGBoost, simple LSTM, and CNN-LSTM without attention become inferior to the proposed approach in all metrics. Moreover, the inference speed test shows that the system meets requirements of real-world applications as transaction inference is completed within 2.3 milliseconds with throughput of more than 21,000 transactions per second.

With the aid of SHAP-based explainability, it becomes clear that the proposed model captures domain-specific fraudulent signals in terms of transaction amount discrepancies, post-transaction balance ratios, transaction time, and inter-transaction intervals. It shows that the model is in sync with existing knowledge about fraud domain and can be readily applied for use within human-centered fraud control procedures stipulated by the financial regulatory framework.

Future research directions would entail investigating the use of federated learning approaches for fraud detection model training on the transaction history datasets from multiple banks simultaneously without disclosing any private customer details, thus addressing the privacy issue and increasing the model's generalizability. In addition, integrating elements of graph neural networks that can learn the graph representations of transactions between accounts is anticipated to aid in detecting collusion among fraudsters, something impossible to achieve with simple transaction history models at the account level. The study of robustness of the proposed model against malicious evasion attacks carried out by experienced fraudsters is another critical research topic. Lastly, the inclusion of multimodal data besides transactional data, such as device fingerprints, geographic location data, and behavioral biometrics, could further enhance detection accuracy.

REFERENCES

1. Y. Liu, Q. Wang, and J. Li, "Deep learning-based fraud detection in financial transactions using attention mechanisms," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 8, pp. 3421–3435, Aug. 2022.
2. R. Mehta, A. Singh, and P. Verma, "Explainable AI for credit card fraud detection: Integrating SHAP with deep neural networks," IEEE Access, vol. 10, pp. 56832–56847, 2022.
3. S. Chen, T. Zhang, and W. Zhou, "Hybrid CNN-LSTM architecture for real-time online banking fraud

ISSN:3048-7722

detection," IEEE Transactions on Information Forensics and Security, vol. 18, pp. 1103–1118, 2023.

4. J. O. Awoyemi, A. O. Adetunmbi, and S. A. Oluwadare, "Credit card fraud detection using machine learning techniques: A comparative analysis," in Proc. IEEE Int. Conf. Computing, Networking and Informatics (ICCNI), Lagos, Nigeria, 2021, pp. 1–9.
5. N. Rtayli and N. Enneya, "Enhanced credit card fraud detection based on SVM-recursive feature elimination and hyper-parameters optimization," Journal of Information Security and Applications, vol. 55, pp. 1–12, 2021.
6. X. Zhang, L. Wu, and M. Chen, "Automated feature engineering for XGBoost-based e-commerce transaction fraud detection," IEEE Transactions on Knowledge and Data Engineering, vol. 35, no. 4, pp. 3889–3902, Apr. 2023.
7. A. Roy and J. Thomas George, "Detecting credit card fraud using deep learning-based LSTM model and GAN for data imbalance handling," in Proc. IEEE Int. Symp. Advanced Networks and Telecommunication Systems (ANTS), Bhopal, India, 2022, pp. 1–6.
8. W. Li, D. Tao, and R. Xu, "A CNN-LSTM hybrid architecture for financial fraud detection with imbalanced sequential data," IEEE Access, vol. 11, pp. 22401–22413, 2023.
9. K. Xie, H. Yang, and L. Zhang, "Self-attention-based sequential model for online banking transaction fraud detection," IEEE Transactions on Computational Social Systems, vol. 11, no. 2, pp. 1544–1556, Apr. 2024.
10. D. Cheng, X. Wang, Y. Zhang, and L. Zhang, "Graph neural network for fraud detection via spatial-temporal attention," IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 8, pp. 3800–3813, Aug. 2022.